Math 711: Lecture of November 29, 2004

We now focus on the situation where the base is a discrete valuation ring $(V,\, tV,\, K)$, so that we are no longer allowing the possibility that $t = 0$.

Recall that $T = V[x_1,\, \dots\, ,x_n]$, a polynomial ring, and that $A = T_{(t,\underline{x})}$, where $\underline{x}$ denotes $x_1,\, \dots\, ,x_n$. Recall also that $S = T[s_1,\, \dots\, ,s_d] \subseteq \widehat{A}$ is a finitely generated T-subalgebra of $\widehat{A}$. Consider the T-algebra map $T[Y_1,\, \dots\, ,Y_d] \twoheadrightarrow S \subseteq \widehat{A}$ such that $y_i \mapsto s_i$, $1 \leq i \leq d$. Let P be the kernel of this map, and let h be the height of P, as before. However, we shall now keep track of a full set of generators of P, which we denote $F_1,\, \dots\, ,F_k$. (Previously, we were working with $F_1,\, \dots\, ,F_h$ such that $PT[Y_1,\, \dots\, ,Y_d]_P$ is generated by $F_1,\, \dots\, ,F_h$.)

Note that t is a prime element of $\widehat{A}$: in fact, $\widehat{A}/t\widehat{A} \cong K[[x_1,\, \dots\, ,x_n]]$. Let $\Lambda = \widehat{A}_{\mathcal{Q}}$ where $\mathcal{Q} = t\widehat{A}$, a DVR with maximal ideal $t\Lambda$. Note that an element of $\widehat{A}$ is divisible by t is and only if that is true in Λ. That is, $t\Lambda \cap \widehat{A} = t\widehat{A}$. We let $\Omega' = \Lambda \otimes_S \Omega_{S/T}$.

Let $\mathcal{J}$ denote the image of the Jacobian matrix $\bigl(\dfrac{\partial F_j}{\partial Y_i}\bigr)$ mod P, which we view as a $d \times k$ matrix over S. Note that the cokernel of $\mathcal{J}$ is $\Omega_{S/T}$. We shall write $\mathcal{J}'$ for $\mathcal{J}$ viewed as a matrix (or linear transformation) over Λ. Thus, the cokernel of $\mathcal{J}'$ may be identified with the module $\Omega' = \Lambda \otimes_S \Omega_{S/T}$.

Note that for any finitely presented module M over any ring R, we can define *Fitting invariants*: the i th Fitting invariant of M is the ideal $I_{d-i}(J)$ where J is the $d \times k$ matrix of the map of free modules in a presentation $R^k \to R^d \twoheadrightarrow M \to 0$.

Fitting's Lemma. *The i th Fitting invariant as defined in the preceding paragraph is independent of the choice of finite presentation of M.*

Proof. To prove this, we first check that given a map $R^d \twoheadrightarrow M$, the i th Fitting invariant is independent of the choice of the finitely many column vectors spanning the kernel. Given two choices, we may compare each with the union. This, it suffices to see that the ideal does not change when one set of relations is included in the other. We may think of one matrix as $d \times k$ and we may assume that it is the submatrix of the other, which is $d \times (k+s)$, formed from the first k columns, while the last s columns are linear combinations of the first k columns. By subtracting linear combinations of the first k columns from the last s (we know that this does not change the ideals of minors) we may assume that the last s columns are all 0, and the result is now clear.

It remains to check independence of the map $R^d \twoheadrightarrow M$, i.e., of the choice of generators for M. Again, we may compare each of two different sets of generators with their union, and so we reduce to the case where one set of generators is included in the other and then, by induction, to the case where there is one additional generator. We may assume that included among the relations is a relation expressing the additional generator, which we number last, as a linear combination of the others. This means that we may assume that the matrix with the additional generators present has a 1 in the last row, which we also assume, by permuting the columns, is in the last column. We can now perform elementary column operations, subtracting multiples of the last column from the others, until the last row consists of all zeros except for its final entry, so that the $(d+1)\times(k+1)$ matrix J_1 that we are considering has the block form $J_1 = \begin{pmatrix} J & C \\ 0 & 1 \end{pmatrix}$, where J is $d \times k$, 0 denotes a row of zeros of length k, C is a $d \times 1$ column, and 1 is the 1×1 identity matrix. Then J gives a matrix for the presentation using the first d generators. It is now straightforward to see that $I_{d+1-i}(J_1) = I_{d-i}(J)$. $\square$

We note that in the Lemma below, the ideal generated by the $h \times h$ minors of $\mathcal{J}$ is the $(d-h)$ th Fitting invariant of the module $\Lambda \otimes_S \Omega$. The number $d-h$ does not depend on the presentation: it is the same as the dimension of $\mathrm{frac}\,(T) \otimes_T S$, which has the presentation $\mathrm{frac}\,(T)[Y_1, \ldots, Y_d]/(F_1, \ldots, F_k)$.

Lemma. *$\mathcal{J}$ has rank h. There is an $h \times h$ minor of $\mathcal{J}$ that is not divisible by t in $\widehat{A}$ (equivalently, in Λ) iff Ω' is torsion-free over Λ. More generally, the minimum order of a size h minor with respect to t is the length over Λ of the torsion submodule of Ω'.*

Proof. Let $\mathcal{F}$ be the fraction field of T, which is also the fraction field of A. Then

$$\mathcal{F} \to \mathcal{F} \otimes_A \widehat{A} = \mathcal{F} \otimes_T \widehat{A}$$

is separable, and, hence, $\mathcal{F} \to \mathcal{F} \otimes_T S$ is separable. Note that $\mathcal{J}$ is also the Jacobian matrix for this extension, and that $\mathcal{F} \otimes_T S$ has dimension $d-h$. If we localize at one nonzero element to make this algebra smooth over $\mathcal{F}$ its dimension does not change, and then the determinantal rank of the Jacobian matrix must be h and, after localization, the size h minors must generate the unit ideal. Since S is a domain, we see that the rank of $\mathcal{J}$ is exactly h.

The remaining statements follow from a general fact about matrices over a DVR, given in the next result. $\square$

Lemma. *Let M be any finitely generated module over a discrete valuation ring $(\Lambda, t\Lambda)$, and let J be the matrix of the map of free modules in a finite presentation of M. Suppose that J has rank h. Then the length of the torsion submodule of M is the same as the minimum order of an $h \times h$ minor of M, which is the same as the order of a generator of $I_h(M)$.*

Proof. The statements are unaffected by performing elementary row and column operations on M. This means that we may assume that J has the block form $\begin{pmatrix} J_0 & 0 \\ 0 & 0 \end{pmatrix}$ where J_0 is a diagonal matrix with diagonal entries $t^{a_1}, t^{a_2}, \ldots, t^{a_h}$. The only nonvanishing $h \times h$ minor is $t^{a_1+\cdots+a_h}$, and $a_1 + \cdots + a_h$ is also the length of the torsion submodule $\Lambda/(t^{a_1}) \oplus \cdots \oplus \Lambda/(t^{a_h})$. $\square$

We want to note the following. Suppose that we have found a size h minor of $\mathcal{J}$ that is not zero. By, renumbering, we may assume that this minor occurs in the upper left hand corner of $\mathcal{J}$, so that it involves the partial derivatives of $F_1, \ldots, F_h$ with respect to $Y_1, \dots, Y_h$. We claim that it is then automatic that $(F_1, \dots, F_h)T[Y_1, \dots, Y_d]_P = PT[Y_1, \dots, Y_d]_P$. Since $T - 0$ does not meet P, we are free to replace T by its fraction field $\mathcal{F}$ in considering this, i.e., we need only show that

$$(F_1, \dots, F_h)\mathcal{F}[Y_1, \dots, Y_d]_P = P\mathcal{F}[Y_1, \dots, Y_d]_P.$$

Let $C = \mathcal{F}[Y_1, \dots, Y_d]/(F_1, \dots, F_h)$. Since there is a maximal minor Δ of the Jacobian matrix that is invertible, $\Omega_{C_\Delta/\mathcal{F}}$ is free of rank $d - h$, and so C_Δ is smooth over $\mathcal{F}$, and therefore regular. If we localize at one additional element g not in P, we may assume that it is a domain. Moreover, it must have dimension $d - h$. Since $C_{\Delta g}$ maps onto $\mathcal{F}[Y_1, \dots, Y_d]_{\Delta g}/P\mathcal{F}[Y_1, \dots, Y_d]_{\Delta g}$, which is also a domain of dimension $d - h$, they are equal, and we must have

$$(F_1, \dots, F_h)T[Y_1, \dots, Y_d]_P = PT[Y_1, \dots, Y_d]_P,$$

as claimed: the two ideals already become equal when we localize at $\Delta g \notin P$ $\square$

We denote by ℓ_S the length of torsion-submodule of $\Lambda \otimes \Omega_{S/T}$, which is also the minimum order with respect to t of any size h minor of J'.

Theorem. *Let $Q = t\widehat{A} \cap S = t\Lambda_0 \cap S$, and let $S' = S[q/t : q \in Q]$. Then $\ell_{S'} \le \ell_S$, and the inequality is strict unless $\ell_S = 0$.*

We put off the proof of this result momentarily. Notice that once we know this we may iterate the process of replacing S by S' at most ℓ_S times until we reach a finitely generated T-subalgebra S^* of $\widehat{A}$ that contains S but such that $\ell_{S^*} = 0$. For this algebra, there will be a size h minor that will not be divisible by t. This enables us to carry through the inductive step in the proof the Artin Approximation Theorem.

However, before we give the proof of this Theorem, we want to observe that the same result also enables us to give a proof for the base case for the induction when V is a DVR and not a field!

To see this, note that in the base case $n = 0$: there are no variables x_i. $T = A = V$, and $\widehat{A} = \widehat{V}$. We may assume that we have $S = V[s_1, \dots, s_d]$ such that $\ell_S = 0$, and we want to map S to V^{h} so that the residue class of every element of S is preserved mod (t^N). Some size h minor of $\mathcal{J}$ is an element of S not divisible by t, that is, it has an inverse in $\widehat{V}$. We may adjoin this inverse to S: constructing the algebra map we need only gets harder. But now we have that S is smooth over V, by the Jacobian criterion for smoothness. Let Q be the ideal $t\widehat{V} \cap S$. By the structure theorem for smooth morphisms, locally S is an étale extension of a polynomial ring, and so S_Q is a pointed étale extension of a localization of a polynomial subring, $V[u_1, \dots, u_r]_{Q'}$. The inclusion $S \subseteq \widehat{V}$ gives us a map $S_{Q'} \to \widehat{V}$, and the u_j have images in $\widehat{V}$. For every j, pick an element of $v_j \in V^{\mathrm{h}}$ such that $u_j \equiv v_j$ mod (t^N). Consider the V-algebra map $V[u_1, \dots, u_r] \to V^{\mathrm{h}}$ sending $u_j \mapsto v_j$ for all j. This extends to a map $V[u_1, \dots, u_r]_{Q'} \to V^{\mathrm{h}}$. Since V^{h} is Henselian, this map extends to the pointed étale extension S_Q, and we have the required map $S \to S_Q \to V^{\mathrm{h}}$. $\square$

This means that after we prove the Theorem stated above, we will have completed the proof of the Artin Approximation Theorem in all cases.